IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Multivariate Active Learning and Adaptive Sampling with Multi-kernel Gaussian Processes

Thien Hoang Nguyen, Nathan Wallace, Nicholas Harrison, and Salah Sukkarieh

***Abstract*—In agriculture, understanding the distribution and relationship between different aspects of the environment is important for minimizing chemical use and reducing environmental impact. Traditionally, it is done by manually collecting samples on the field and then sending them to a laboratory for analysis. This is not only labor-intensive and costly, but the results will still be outdated. There is thus a growing interest in developing robotic systems to map these variables and uncover their correlations in real time. However, existing learning and sampling methods only focus on one quantity of interest (QoI) or make assumptions that might lead to sub-optimal results when there are multiple QoIs. In this work, we propose a multivariate active transfer learning and intelligent adaptive sampling system that can simultaneously learn the most accurate models for multiple QoIs as well as the relationship between them, and leverage that knowledge to select the next best locations to sample. Performance benchmarking against existing methods shows that QoIs are mapped more accurately, complex correlations between QoIs are identified more precisely, and travel routes are planned more efficiently.**

***Index Terms*—Adaptive Sampling, Active Transfer Learning, Multi-kernel Gaussian Process, Precision Agriculture**

I. Introduction

GAINING real-time knowledge of the environment is a fundamental problem for many applications. In farming, for example, knowing the levels of macronutrients such as NPK (nitrogen, phosphorus, potassium), pH, water, and crop health is critical. With accurate and timely information, farmers can determine the amount of fertilizers and chemicals required to increase yield in a sustainable manner. However, the traditional way of sampling the field is a manual and labor-intensive process. This hampers our ability to fully understand the interdependencies among these factors and limits our capacity to optimize resources in farming practices [1].

To address this problem, autonomous robots play a decisive role due to the power of autonomy and mass deployment. Across different fields, many robotic systems are being deployed to perform tasks that are dull, dirty and even dangerous. In precision agriculture (PA), robots can be deployed to rapidly collect samples to measure the macronutrient composition and water characteristics at various depths below the surface and various locations across the field, at a pace that far surpasses humans [2]. Hence, robots will be key to enabling automated and large-scale PA of the future.

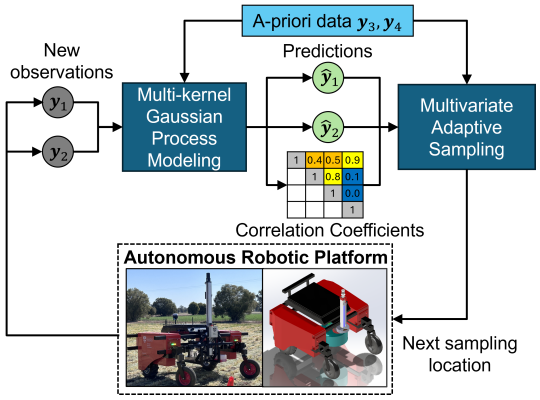

Fig. 1: Overview of the proposed system in the case with $q = 2$ QoIs and $p = 2$ PQ. The system operates in two phases: learning (MKGP module) and sampling (MVAS module).

However, the full potential of agricultural robots can only be unlocked with novel algorithms that can make the most optimal decisions given the situation. Firstly, using in-situ measurements and prior knowledge obtained from past missions or other sources of information, the algorithm should incorporate spatial process modeling, making inferences and predicting at unsampled locations. Secondly, the algorithm should provide information-guided sampling strategies, i.e. it should prioritize effective samples which will improve the model's accuracy and reduce the overall uncertainty. Lastly, the algorithm should consider the robot's budgetary constraint since it is a major concern for real-life operation.

In this context, we propose a system that combines the power of active transfer learning and intelligent adaptive sampling for environmental monitoring applications. Our main contributions include:

- *MKGP*: a multi-kernel, multi-output modeling method that can learn the spatial models of and correlation among many quantities of interest (QoIs) and prior data;
- *MVAS*: a multivariate adaptive sampling method that leverages the model from MKGP, a-priori data and travel cost to select the most effective sampling locations.
- A thorough evaluation process for each component of the system with synthetic and real-world datasets.

The rest of the paper is structured as follows: First, the related works are reviewed in Sect. II. Next, we present our system in detail in Sect. III. Then, Sect. IV evaluates our system in simulations and Sect. V concludes our work.

Manuscript received: October 3, 2024; Revised March 18, 2025; Accepted May 9, 2025.

This paper was recommended for publication by Editor Hyungpil Moon upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by the Vonwiller Foundation and the University of Sydney's Digital Sciences Initiative.

The authors are with the Australian Centre for Robotics, University of Sydney, NSW, AU (e-mail: {thienhoang.nguyen, nathan.wallace, nicholas.harrison, salah.sukkarieh}@sydney.edu.au).

Digital Object Identifier (DOI): see top of this page.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

II. Literature Review

In many applications where a robot is sent to locations of interest to gather environmental information, an appealing modeling technique is the Gaussian Process (GP) [3]–[6]. From only sparse data, GPs can estimate both values and uncertainties of spatially-dependent quantities. However, most previous works rely on single-output GPs (SOGP), i.e., GPs that are used to predict a single output based on the inputs. For PA, we need to explore not only the properties of the environment but also the complex correlations between them. Hence, there is a need to develop multivariate learning and sampling methods for autonomous robots.

Recent works have employed multi-output GPs (MOGP) for model learning in robotics. In [7], using MOGP was shown to improve the accuracy in learning a robot's inverse dynamics compared to using multiple SOGPs. In [8], the authors presented a MOGP system for occupancy grid map generation using camera and lidar data from multiple robots. In [9], we introduced an MOGP-based modeling and optimization-based sample selection method for information-gathering applications with non-collocated data. However, these works are single-kernel MOGP (SKGP for short) methods, which means they assume the same kernel type and the same hyperparameters can be used to model all tasks. While this assumption can reduce the complexity of the optimization problem, in real applications different kernel types or at least different hyperparameters may be required to model each task accurately [10], [11]. Hence, developing modeling methods using multi-kernel MOGP (MKGP for short) to address this limitation is the logical next step.

One major challenge for MOGP is the high computational cost of inverting large covariance matrices, which grows with the number of tasks and data points [12]. To tackle this issue, various solutions have been proposed. For example, [13] combines deep neural networks with sparse GPs to provide fast, reliable uncertainty estimates. However, this method relies on having a sufficiently large dataset for both training and inference. [14] introduces an intrinsically sparse covariance function that reduces inversion complexity but sacrifices some uncertainty quantification. In contrast, MKGP prioritizes prediction accuracy at the cost of higher regression complexity, which is more aligned with PA scenarios where data is sparse, the ability to collect large quantities of data is limited, and prediction accuracy is paramount.

Targeting applications with collocated prior data, which are increasingly common due to satellite imagery, sensor networks, and geostatistical surveys, our system improves upon previous works in the following ways: Firstly, compared to SOGP and SKGP methods [3]–[5], [9], our approach allows the most suited kernel type to be assigned to each QoI. Secondly, our method explicitly leverages prior data in the sample selection phase, while [5], [9] only include it implicitly within the model. Thirdly, our system incorporates a multivariate objective function to plan a more effective route for all QoIs, while past approaches only consider one QoI at a time. As a result, we observe noticeable enhancements in performance compared to existing methods, which will be discussed in Sect. IV. Table I summarizes the comparison between our MKGP and related works.

Method	Learning		Sampling	
	Output	Kernel	Location	Cost
mGP [4]	Single	Single	✓	-
MI [3]	Single	Single	✓	✓
SOGP [5]	Single	Single	✓	✓
SKGP [9]	Multi	Single	✓	-
Ours	Multi	Multi	✓	✓

TABLE I: Comparision of our method and related works.

Recent advances in adaptive sampling have extended traditional myopic and lookahead strategies to multivariate scenarios, where multiple QoIs are considered simultaneously [15]. These approaches typically balance exploration and exploitation by jointly assessing the predictive uncertainty and expected improvement across all outputs. [16] presents a highly scalable sampling-based planning method for multi-robot information gathering missions in complex environments. However, in these works the performance guarantee is on uncertainty reduction rather than modeling accuracy. In our system, MVAS is designed specifically to improve the modeling accuracy of MKGP, by explicitly incorporating the resource constraint and parameter-wise correlations in our formulation. Our work adopts an information-based heuristic that quantifies the expected improvement in mapping accuracy and reduction in model uncertainty at candidate sampling points, while incorporating travel cost into the decision process. By doing so, our method provides a computationally efficient and principled way to guide sample selection in multivariate settings without the added complexity of full multivariate lookahead strategies.

III. Proposed System

A. Problem Formulation

Suppose we have a 2D environment with a grid map that represents the traversability information. A robot is deployed to measure $q \in \mathbb{N}$ scalar QoIs at each sample location. Additionally, we have data for $p \in \mathbb{N}$ scalar quantities (referred to as "prior quantities" - PQ) available for all sample locations but their correlations with the QoIs are unknown. Let $n = q+p$ be the total number of quantities. The problem is mapping the distribution of all q QoIs and the correlation among all n quantities as accurately and efficiently as possible. Fig. 1 shows an overview of our system and its main components.

B. Multi-kernel Gaussian Process modeling (MKGP)

1) Modeling: Let $\mathbf{x} \in \mathbb{R}^2$ be the sampling location and $\mathbf{y}(\mathbf{x}) = [y_1(\mathbf{x}), \dots, y_n(\mathbf{x})]^\top \in \mathbb{R}^n$ be the observed data at $\mathbf{x}$. The first q and the last p elements of $\mathbf{y}(\mathbf{x})$ are in-situ measurements for the QoIs and a-priori data for PQ, respectively. Let $\mathbf{S} = \{\mathbf{X}, \mathbf{Y}\}$ be the training dataset with m data points where $\mathbf{X} = \{\mathbf{x}_j \mid j = 1, \dots, m\}$ and $\mathbf{Y} = \{y_i^j \mid i = 1, \dots, n, j = 1, \dots, m\}$. In SOGP, the aim is to learn n independent regression models, or tasks. In MKGP, we consider $\mathbf{Y}$ to be drawn from an n-variate GP with

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}_{\mathbf{Y}}$. $\boldsymbol{\Sigma}_{\mathbf{Y}}$ describes the coupling between the outputs according to the distance between the inputs and the correlation between the tasks. Generally, $\boldsymbol{\Sigma}_{\mathbf{Y}} \in \mathbb{R}^{mn \times mn}$ takes the form [17]:

$$\boldsymbol{\Sigma}_{\mathbf{Y}} = \begin{bmatrix} \mathbf{C} & \boldsymbol{\Sigma}(\mathbf{x}_1, \mathbf{x}_2) & \dots & \boldsymbol{\Sigma}(\mathbf{x}_1, \mathbf{x}_m) \\ \boldsymbol{\Sigma}(\mathbf{x}_1, \mathbf{x}_2) & \mathbf{C} & \dots & \boldsymbol{\Sigma}(\mathbf{x}_2, \mathbf{x}_m) \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{\Sigma}(\mathbf{x}_1, \mathbf{x}_m) & \boldsymbol{\Sigma}(\mathbf{x}_2, \mathbf{x}_m) & \dots & \mathbf{C} \end{bmatrix}, \quad (1)$$

where:

$$\mathbf{C} \coloneqq \boldsymbol{\Sigma}(\mathbf{x}_a, \mathbf{x}_a) = \begin{bmatrix} \sigma_1^2 & \sigma_{1,2} & \cdots & \sigma_{1,n} \\ & \sigma_2^2 & \cdots & \sigma_{2,n} \\ & & \ddots & \vdots \\ & & & \sigma_n^2 \end{bmatrix} \quad (2)$$

is the $n \times n$ cross-covariance matrix between the tasks, $\sigma_{i,j}$ is cross-covariance between the i-th and j-th tasks, $\boldsymbol{\Sigma}(\mathbf{x}_a, \mathbf{x}_b)$, $1 \leq a, b \leq m$, is the $n \times n$ covariance matrix between the outputs at locations $\mathbf{x}_a$ and $\mathbf{x}_b$. This form stacks the covariance matrices per input combination and not per output, which we leverage to derive the general formula of $\boldsymbol{\Sigma}(\mathbf{x}_a, \mathbf{x}_b)$ from $\mathbf{C}$. In [11], it was shown that $\boldsymbol{\Sigma}(\mathbf{x}_a, \mathbf{x}_b)$ can be constructed from base kernels that represent spatial and inter-task correlations. However, while [11] assumes unit variances for each kernel, we leverage the fact that the signal variances can be explained fully by $\mathbf{C}$ to define:

$$\boldsymbol{\Sigma}(\mathbf{x}_a, \mathbf{x}_b) = \mathbf{C}(\boldsymbol{\theta}_C) \odot \mathbf{K}(\mathbf{x}_a, \mathbf{x}_b, \boldsymbol{\theta}_K), \quad (3)$$

where $\odot$ is the Hadamard product (element-wise product), $\mathbf{K}$ is a $n \times n$ multi-kernel covariance matrix, $\boldsymbol{\theta}_C$ and $\boldsymbol{\theta}_K$ are the vectors containing the hyperparameters for $\mathbf{C}$ and $\mathbf{K}$, respectively. Next, we outline how to construct $\mathbf{K}$ and $\mathbf{C}$.

$\mathbf{K}$ consists of auto covariance terms K_{ii} ($i \in [1, n]$) and cross-covariance terms K_{ij} ($i \neq j$), all of which need to be defined. The auto covariance terms are selected first from existing kernel choices to best model task-specific spatial behavior. For example, the Matérn covariance function is a single-quantity covariance function commonly used in probability theory, machine learning and spatial statistics [18]. The Matérn 3/2 covariance function is defined as:

$$M(\mathbf{x}_a, \mathbf{x}_b, \theta_i) = \left(1 + \sqrt{3}d/\theta_i\right) \exp\left(-\sqrt{3}d/\theta_i\right), \quad (4)$$

where $d = \|\mathbf{x}_a - \mathbf{x}_b\|$ is the geodesic distance between sample locations, θ_i is the corresponding characteristic length-scales. If we choose K_{ii} and K_{jj} as Matérn kernels with different hyperparameters $\theta_{K,i}$ and $\theta_{K,j}$, the cross-covariance function can be analytically computed as [11]:

$$K_{ij} = \theta_{ij}\left(\theta_{K,i} e^{-\sqrt{3}d/\theta_{K,i}} - \theta_{K,j} e^{-\sqrt{3}d/\theta_{K,j}}\right), \quad (5)$$

where $\theta_{ij} = 2\sqrt{\theta_{K,i}\theta_{K,j}}/(\theta_{K,i}^2 - \theta_{K,j}^2)$. We note that when $\theta_{K,i} \approx \theta_{K,j}$ Eq. (5) is simplified into $K_{ij} = e^{-\sqrt{3}d/\theta_{K,i}}$. Other kernel combinations as well as the process to obtain any arbitrary combination can be found in [11].

We then follow the free-form parameterization [19] to construct $\mathbf{C} = \mathbf{L}\mathbf{L}^\top$, where $\mathbf{L} \in \mathbb{R}^{n \times n}$ is a lower triangular matrix. As such, the non-zero elements of $\mathbf{L}$ will be the cross-correlation hyperparameters, i.e., $\boldsymbol{\theta}_C = [L_{11}, L_{21}, L_{22}, \cdots, L_{nn}]^\top \in \mathbb{R}^N$ where $N = n(n+1)/2$.

With the proposed approach, $\mathbf{K}$ and $\mathbf{C}$ are guaranteed to be positive semi-definite (PSD), which guarantees $\boldsymbol{\Sigma}_{\mathbf{Y}}$ is also PSD and the resulting Gaussian distributions are well-defined. Moreover, in [20], it was shown that if the individual processes are GPs, then the convolved process is also a valid GP. As such, existing theoretical results for GPs still apply to MKGPs. In particular: (1) while global convergence is not guaranteed, increasing the size of the training set typically improves hyperparameter estimation and modelling accuracy [18], and (2) the approximation error depends on how well the model aligns with the underlying assumptions (e.g., noise characteristics, likelihood choice) [21].

2) Hyperparameter Optimization and Correlation Estimation: The hyperparameters, which include $\boldsymbol{\theta}_C$ and $\boldsymbol{\theta}_K$, are estimated by the restricted maximum likelihood estimation (RMLE). RMLE has been shown to provide less biased estimation results compared to MLE even with a low number of samples [17]. RMLE is written as:

$$\text{RMLE} = \ln|\boldsymbol{\Sigma}_{\mathbf{Y}}| + \ln|\mathbf{F}^\top \boldsymbol{\Sigma}_{\mathbf{Y}} \mathbf{F}| + (\mathbf{Y} - \mathbf{F}\hat{\boldsymbol{\mu}})^\top \boldsymbol{\Sigma}_{\mathbf{Y}}^{-1} (\mathbf{Y} - \mathbf{F}\hat{\boldsymbol{\mu}}) \quad (6)$$

where $\mathbf{F} = \mathbf{1}_m \otimes \mathbf{I}_n$, $\mathbf{1}_m$ is a $m \times 1$ vector of ones, $\mathbf{I}_n$ is an $n \times n$ identity matrix, $\otimes$ is the Kronecker product, and $\hat{\boldsymbol{\mu}}$ is the generalized least-squared estimator [22]:

$$\hat{\boldsymbol{\mu}} = (\mathbf{F}^\top \boldsymbol{\Sigma}_{\mathbf{Y}}^{-1} \mathbf{F})^{-1} \mathbf{F}^\top \boldsymbol{\Sigma}_{\mathbf{Y}}^{-1} \mathbf{Y}. \quad (7)$$

Unlike traditional MLE, which can often be optimized directly using gradient descent, the RMLE objective function in Eq. (6) may exhibit multiple local maxima. To address this, we first obtain initial parameter estimates using a global optimization method—specifically, a particle swarm optimization routine—and then refine the estimates using gradient descent. After optimization, we need to normalize the estimated $\hat{\mathbf{C}}$. The reason is that as the free-form parameterization in [19] does not force the diagonal elements of $\hat{\mathbf{C}}$ to be 1, the values of the correlations will not be limited to $[-1, 1]$, which makes interpreting the inter-task similarities less straightforward. To address this issue, we follow [12] to obtain the normalized matrix $\hat{\mathbf{C}}'$. $\hat{\mathbf{C}}'$ will have 1 as diagonal elements, and each element at the a-th row and b-th column ($b < a$) is computed as $\hat{C}'_{ab} = \hat{C}_{ab} / \sqrt{\sum_{\delta=\delta_1}^{\delta_2} \theta_{C,\delta}^2}$, where $\delta_1 = a(a-1)/2 + 1$ and $\delta_2 = a(a+1)/2$. We then define the correlation coefficient $r_{ab} = (\hat{C}'_{ab})^2$ to give a score from 0 to 1 that represents how much one task is explained by the other. The closer r_{ab} is to 1, the more correlated task a and b are. Conversely, when r_{ab} is close to 0, the tasks are more likely to be independent.

3) Prediction: Once $\boldsymbol{\theta}_C$ and $\boldsymbol{\theta}_K$ are estimated, the multivariate best linear unbiased prediction is computed as [17]:

$$\hat{\mathbf{y}} = \hat{\boldsymbol{\mu}} + \boldsymbol{\Sigma}_{\mathbf{T}} \boldsymbol{\Sigma}_{\mathbf{Y}}^{-1} (\mathbf{Y} - \mathbf{F}\hat{\boldsymbol{\mu}}), \quad (8)$$

where $\boldsymbol{\Sigma}_{\mathbf{T}}$ is the covariance matrix between the testing point $\mathbf{x}^*$, whose outputs $\hat{\mathbf{y}}$ are to be predicted, and each of the training points in $\mathbf{X}$, whose outputs $\mathbf{Y}$ are known. Lastly,

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

following [17], the $n \times n$ matrix of estimated mean squared prediction error (MSPE) is computed as:

$$\boldsymbol{\varepsilon}(\mathbf{x}^*) = \mathbf{C} - \boldsymbol{\Sigma}_{\mathbf{T}}\boldsymbol{\Sigma}_{\mathbf{Y}}^{-1}\boldsymbol{\Sigma}_{\mathbf{T}}^{\top} + \mathbf{U}(\mathbf{F}^{\top}\boldsymbol{\Sigma}_{\mathbf{Y}}^{-1}\mathbf{F})^{-1}\mathbf{U}^{\top}, \quad (9)$$

where $\mathbf{U} = \mathbf{I}_n - \boldsymbol{\Sigma}_{\mathbf{T}}\boldsymbol{\Sigma}_{\mathbf{Y}}^{-1}\mathbf{F}$.

C. Multivariate Adaptive Sampling (MVAS)

Given a training set $\mathbf{S}$ with m data points, the goal of our Multivariate Adaptive Sampling (MVAS) algorithm is to find the next sampling location $\mathbf{x}_{m+1}$ with the following criteria: (1) the new observation at $\mathbf{x}_{m+1}$ should improve the model accuracy and reduce the uncertainty for all QoIs, (2) a-priori data should be leveraged if deemed appropriate by the inter-task correlations, and (3) the algorithm should aim to minimize the travel cost. We propose the following MVAS optimization problem that satisfies these criteria as:

$$\max_{\mathbf{x}\in\boldsymbol{\mathcal{X}}} \frac{\mathcal{I}(\mathbf{x})}{\Lambda(\mathbf{x}, \mathbf{x}_m)}, \quad (10)$$

where $\boldsymbol{\mathcal{X}}$ is the set of unvisited locations and $\Lambda(\mathbf{x}, \mathbf{x}_m) = \|\mathbf{x} - \mathbf{x}_m\|/v$ is the travel cost from $\mathbf{x}_m$ to $\mathbf{x}$, with v as the robot velocity. Note that our method can be generalized to any arbitrary travel cost functions $\Lambda(\cdot)$. Based on Expected Improvement for Global Fit [23], $\mathcal{I}(\cdot)$ is a function that estimates the information gained by visiting $\mathbf{x}$ as:

$$\mathcal{I}(\mathbf{x}) = \overbrace{\|[\hat{\mathbf{y}}(\mathbf{x}) - \mathbf{y}(\mathbf{x}_*)]_q\|^2 + \alpha \sum_{i=1}^{q} \varepsilon_{ii}(\mathbf{x})}^{\text{First}} + \overbrace{\beta \sum_{i=1}^{p} \bar{r}_i \varphi_i^2(\mathbf{x})}^{\text{Second}}, \quad (11)$$

where $\mathbf{x}_*$ is the closest sampled point to the candidate point $\mathbf{x}$ in terms of Euclidean distance, $[.]_q$ extracts the first q elements of the input vector, $\hat{\mathbf{y}}(\mathbf{x})$ and $\boldsymbol{\varepsilon}(\mathbf{x})$ are respectively the predicted value and MSPE computed in Sect. III-B3, $\bar{r}_i = \max_{j\in[1,q]} \hat{C}'_{q+i,j}$ is the maximum of all the inter-task correlation coefficients between PQ_i and the QoIs, $\varphi_i = |\hat{y}_i - \check{y}_i|$ is the error between the predicted ($\hat{\cdot}$) and a-priori ($\check{\cdot}$) values for PQ_i, and α and β are user-defined parameters to change the balance of the optimization. φ_i^2 is used instead of φ_i to make the units in Eq. (11) consistent.

The MVAS problem (10) essentially maximizes the information gain efficiency when visiting each new location. It can be interpreted by understanding how the components are designed. First, the information gain function in Eq. (11) aims to satisfy criteria (1) and (2) as follows. The First term can improve the mapping accuracy while the Second term encodes the information that a-priori data can provide. If $\bar{r}_i \approx 1$, there is a strong correlation between PQ_i and at least one QoI. As such, measurements that can quickly improve the internal model for PQ_i (via φ_i^2) may also quickly improve the model for the associated QoI. Hence, the algorithm should also favor locations that can significantly reduce the errors of the internal model for PQ_i. Conversely, when $\bar{r}_i \approx 0$, meaning PQ_i is uncorrelated with any QoI, the Second term will be insignificant and can be ignored. Finally, we divide the estimated information gain by the travel cost to satisfy criterion (3): the algorithm should favor locations that might be less informative but are much cheaper to travel to, which is more cost-effective from an operational standpoint. While MVAS is not derived directly from Shannon's information theory, they are conceptually aligned. Shannon's entropy measures uncertainty, and information gain is the reduction in that uncertainty upon acquiring new data. Our information-based heuristic method draws on the same underlying idea of quantifying and reducing uncertainty similar to Shannon's information theory, but they are not formally equivalent.

IV. Simulation Results

In this section, we evaluate our system in simulation, focusing on key metrics related to a robotic problem. Specifically, we show that compared to existing approaches: 1) MKGP can learn more accurate models for the QoIs, 2) MVAS can plan more effective and efficient sampling routes, and 3) our system, particularly MKGP, is more computationally demanding. For more in-depth comparisons between SOGP, SKGP and MKGP techniques, interested readers are referred to [12], [17], [19].

A. Synthetic Data

1) Simulation Setup: We consider a practical scenario in agriculture where we wish to use a robotic solution to map the distribution of Nitrogen (N) and Phosphorous (P), two important macronutrients for plant and fruit growth, on a farm field. The distribution of Potassium (K), another macronutrient, can be obtained from a recently conducted geochemical survey such as [24], and the distribution of pH can be obtained from a sensor network available at the field. We simulate a $50\,\text{m} \times 50\,\text{m}$ field with grid size 1, i.e. each measurement represents a $1\,\text{m} \times 1\,\text{m}$ area, and generate $q = 2$ QoIs ("N" and "P") and $p = 2$ PQ ("K" and "pH") as follows. Firstly, each QoI_i ($i \in [1, q]$) is generated using the method in [25], which is a well-established method in the literature that ensures the synthetic data exhibits spatial correlations and statistical properties similar to real-world data. The observation values for the QoIs are in $[0, 100]$ range, which are corrupted by Gaussian noise $\epsilon \sim \mathcal{N}(0, 10)$. Then, each PQ_i ($i \in [1, p]$) is generated from one of three methods:

- **High correlation (H)**: PQ_i is QoI_i multiplied by a random positive value, i.e. $\text{PQ}_i = \eta_i \text{QoI}_i$, $0 < \eta_i \leq 1$,
- **Medium correlation (M)**: PQ_i is QoI_i added with large Gaussian noise, i.e. $\text{PQ}_i = \text{QoI}_i + \eta_i$, $\eta_i \sim \mathcal{N}(0, 30)$,
- **Low correlation (L)**: the map of PQ is a completely new map, independent of any QoI.

In an agricultural context, these correlation levels may indicate normal plant growth stages (high/medium correlation between N-K or P-K), or potential issues such as nutrient leaching or soil acidification (high/medium correlation between N-pH or P-pH). Next, we separately evaluate each component of our system to better understand their individual performance.

2) Modeling Accuracy: We compare MKGP against SOGP and SKGP methods using the same training set $\mathbf{S}$ obtained from a coverage path with $m = 20$ data points. The Matérn 3/2 kernel type was chosen for all learning methods. The root-mean-square error (RMSE) is calculated as $\text{RMSE}_i =$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

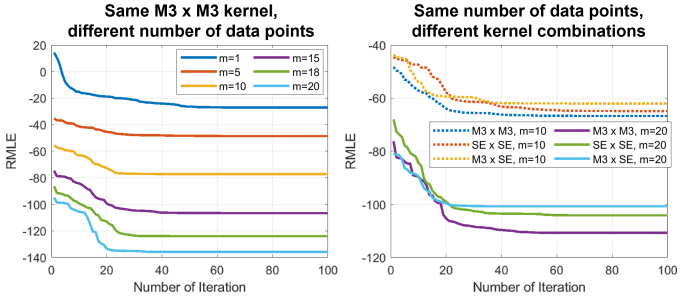

Fig. 2: Training RMLE of the proposed MKGP method with increasing number of data points m (left) and different kernel combinations (right). M3 and SE refer to Matérn 3/2 and Squared Exponential kernel types, respectively.

	Method	Configuration of QoI_i-PQ_i					
		LL	LM	LH	MM	MH	HH
Learning	SOGP	14.02	15.16	13.85	15.75	11.61	17.19
	SKGP	14.84	15.47	15.10	15.20	12.36	17.01
	MKGP	**12.36**	**14.79**	**12.03**	**15.10**	**11.50**	**12.99**
Sampling	RS	16.58	16.41	17.16	15.12	16.62	13.67
	CS	16.76	16.31	18.09	15.68	18.82	12.82
	EIGF	16.64	15.92	16.51	14.77	16.51	11.98
	MVAS	**15.20**	**14.78**	**13.75**	**13.34**	**14.29**	**11.51**

TABLE II: Simulation results. We report the average RMSE of the predicted map over all QoIs. Configuration refers to the correlation level between each QoI_i-PQ_i pair, e.g."LM" means N (QoI_1) has a low correlation with K (PQ_1) and P (QoI_2) has a medium correlation with pH (PQ_2). The best and second best results are in **bold** and blue, respectively.

$\sqrt{\frac{1}{M}\sum_{j=1}^{M}(\tilde{y}_i(\mathbf{x}_j) - \hat{y}_i(\mathbf{x}_j))^2}$ where M is the number of test points (i.e., number of grid cells in the map), $\tilde{y}_i(\mathbf{x}_j)$ and $\hat{y}_i(\mathbf{x}_j)$ are the true and predicted values at location $\mathbf{x}_j$. RMSE is used to evaluate the mapping accuracy for each QoI_i, $i \in [1, p]$. Each experiment is repeated 10 times with different QoIs, then the average RMSE over all RMSE_i ($\overline{\text{RMSE}}$) is reported in Table II. This was enough runs to ensure the statistical significance of our findings, i.e. a 5% p-value.

Fig. 2 illustrates the RMLE of MKGP during one simulation. As can be seen, our model can converge after 20-50 iterations in most cases. We can use the RMLE value as the training metric, where the lower RMLE is the more accurate the model becomes. Based on this metric, it can be said that the trained model becomes more accurate with more training data. Furthermore, among the kernel combination choices, Matérn 3/2 $\times$ Matérn 3/2 has the fastest learning rate and lowest RMLE after training, which might indicate that it is the most suitable for our application. These observations demonstrate that our modeling method can converge and has behaviours within established knowledge about GPs [11], [18].

As can be seen from Table II, the proposed MKGP outperforms other methods in all configurations. Furthermore, the results of SOGP are generally better than SKGP in most cases, which shows that assigning a different kernel for each QoI is generally better than having a single kernel for all QoIs. However, by also learning the correlations between data, our method can improve the modeling accuracy noticeably and

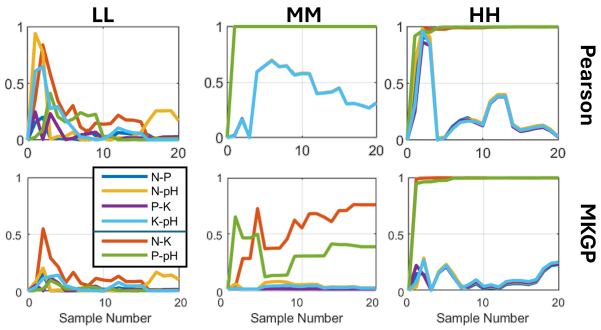

Fig. 3: Comparision between MKGP and Pearson methods in estimating the correlation coefficients r_{ab} for linear relationship. LL/MM/HH refers to the correlation configuration of both N-K and P-pH pairs. The other pairs are uncorrelated.

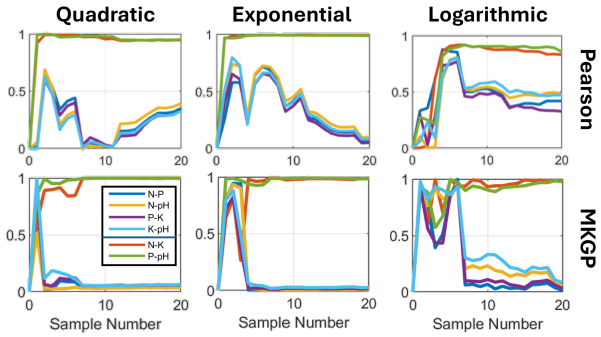

Fig. 4: Comparision between MKGP and Pearson methods in estimating the correlation coefficients r_{ab} for three types of nonlinear relationship. N-K and P-pH share the same type of relationship in each case, while other pairs are uncorrelated.

was able to surpass SOGP in all cases.

3) Correlation Estimation Accuracy: We evaluate the effectiveness of our method in estimating the correlation between quantities by comparing the correlation coefficient r_{ab} estimated by MKGP and the bivariate Pearson correlation, which is commonly used to measure the correlations among pairs of variables. We tested the methods against 3 types of linear relationships, as listed in Sect. IV-A, and 3 types of nonlinear relationships as follows:

- **Quadratic correlation:** $\text{PQ}_i = 0.5\text{QoI}_i^2 + 0.1\text{QoI}_i$, as QoI_i increases, PQ_i increases at an accelerating rate,
- **Exponential correlation:** $\text{PQ}_i = 0.3e^{\text{QoI}_i}$, as QoI_i increases, PQ_i increases exponentially,
- **Logarithmic correlation:** $\text{PQ}_i = \log(\text{QoI}_i)$, as QoI_i increases, PQ_i increases at a decreasing rate.

Fig. 3 and Fig. 4 illustrate the results in linear and nonlinear cases, respectively. While both methods can correctly identify the correlation, MKGP can do so with only one sample whilst the Pearson method requires at least 3 samples to reach the same outcome. Furthermore, for the uncorrelated pairs of quantities, the results from the Pearson method fluctuate much more and stabilize at a much slower rate than the proposed method. In some challenging cases such as linear HH in Fig. 3 and nonlinear Logarithmic in Fig. 4), the Pearson method

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

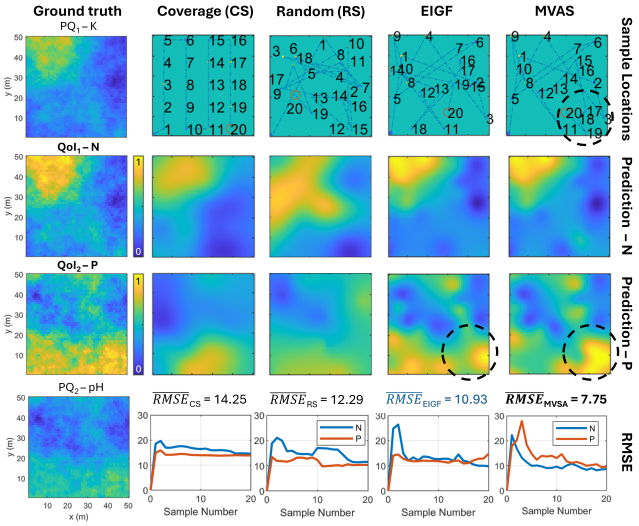

Fig. 5: Comparison between different sampling methods in one simulation case with linear HH configuration. Values have been normalized within each map. Cooler and warmer colours indicate lower and higher abundances, respectively. The circles highlight one area where MVAS was able to take advantage of a-priori data to improve the performance compared to EIGF.

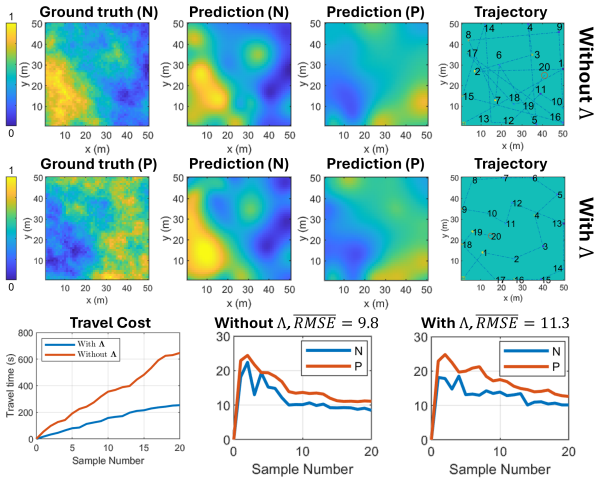

Fig. 6: Example of mapping accuracy and travel cost when $\Lambda(\cdot)$ is and is not taken into account.

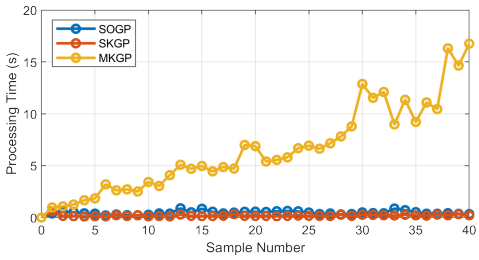

Fig. 7: Comparision of computation time between different learning methods. The results are averaged over 10 runs.

could not reach a definite conclusion for the uncorrelated pairs, while the proposed MKGP can accurately identify the corresponding relationship, or lack thereof. As such, MKGP can more confidently ascertain the correlation between any pair of quantities, for both linear and nonlinear cases.

4) Sampling Performance: To evaluate the sampling module, we pair MKGP with one of the following sampling methods: random sampler (RS), which selects the next sample location randomly; coverage sampler (CS), which selects the next sample location following a grid pattern; Expected Improvement for Global Fit (EIGF), which is the sample selection method used in [9] extended for multiple QoIs by adding individual objectives together; and MVAS. The coefficients in Eq. (11) are empirically chosen as $\alpha = 100$, $\beta = 100$. MKGP is initialized with $n = 4$ data points obtained near the starting point ($[0, 0]$), then we let the algorithms run until a budget of $M = 20$ samples is spent. As the other methods did not take the travel cost into account, we removed $\Lambda(\cdot)$ from the MVAS cost function (10) to make the comparison fair.

Table II shows the RMSE results for the sampling methods, while Fig. 5 demonstrates an example of the difference between them. From Table II, it is obvious that adaptive sampling-based methods are much more effective than a coverage path or randomization. Furthermore, the maps provided by adaptive sampling methods are more representative of the underlying fields, as shown in Fig. 5. However, our method can exploit a-priori data to choose the next sampling location more effectively than EIGF. As a result, MVAS allows the robot to focus on areas with higher potential for improvement, which leads to higher map fidelity (as seen in Fig. 5) and overall better accuracy (as shown in Table II).

5) Travel and Computation Costs: We compare the performance of our full system with and without taking the travel cost into account, i.e. by removing $\Lambda(\cdot)$ from Eq. (11). We run both configurations in 10 simulations with random initialization conditions and the robot velocity set at $v = 1$ m/s. Overall, without penalizing for the travel cost, the mapping accuracy of MVAS can increase by 5% on average. This is expected, as the performance should be better if information gain is the only target for optimization. However, when $\Lambda(\cdot)$ is included, the travel cost is reduced by 41%, which is a significant improvement. Fig. 6 shows an example of the difference of the two cases. It can be seen that, with some small variations, the sampling routes go through the same locations in both cases. However, MVAS simply selects less informative but closer locations first before visiting more informative but further away locations later on. Hence, by optimizing for the right objective, MVAS can become much more cost-effective at the expense of slightly lower accuracy. Thus, the proposed MVAS is suitable for real-world applications, where cost is a major operational factor.

All simulations were run on a laptop with Intel i7-13700H CPU, 32GB RAM, and without any GPU acceleration. Fig. 7 illustrates the average processing time over 10 runs of different learning methods, including learning hyperparameters and inference. It is clear that MKGP is more computationally heavy than other approaches and suffers from a scaling issue. As such, we note some limitations of our system that lead to high computation costs. First is the operations involving the large matrix $\mathbf{\Sigma_Y}$, in particular the inversion in Eq. (6).

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

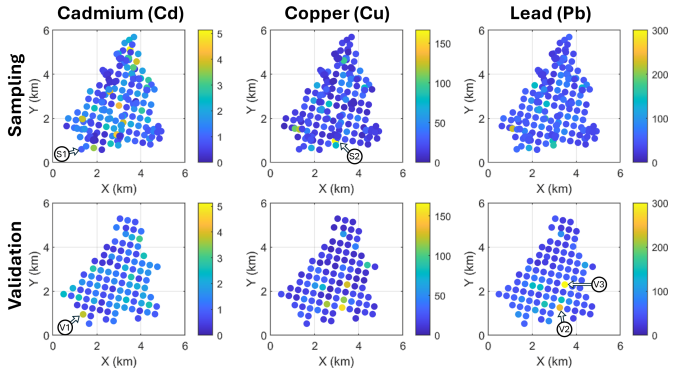

Fig. 8: True spatial distribution maps of Cd, Cu, and Pb in the Jura dataset. The robot's starting location is labeled as (S1). The unit for all maps is parts per million (ppm).

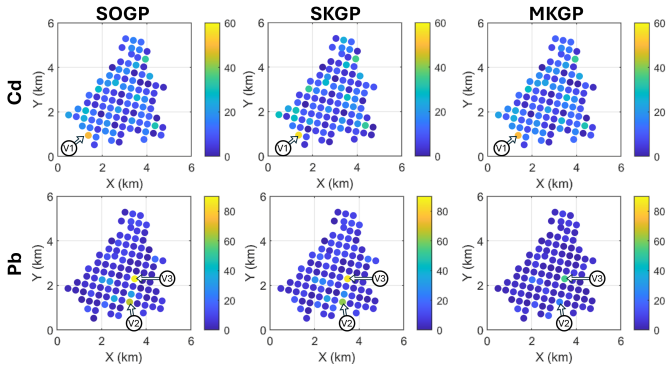

Fig. 9: Comparison of the percent errors of different modelling methods on Jura dataset, with all sampling data processed.

Method	Mean PE (%)		Statistics for PE (%)			Correlation	
	Cd	Pb	σ	Max	Min	Cd-Cu	Pb-Cu
SOGP	15.07	14.27	11.94	87.22	0.08	**0.13**	0.61
SKGP	16.41	14.72	10.97	84.54	0.08	0.14	0.60
MKGP	**14.55**	**13.25**	**4.95**	**53.41**	**0.06**	0.18	**0.79**

TABLE III: Experimental results for different learning methods on Jura dataset, with all sampling data processed. The best and second best results are in **bold** and blue, respectively.

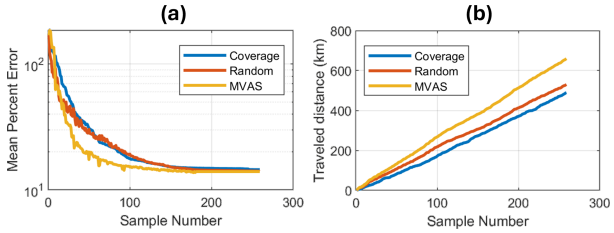

Fig. 10: Comparison of different sample selection methods on Jura dataset. a) Mean percent error averaged over both Cd and Pb. b) Total traveled distance by the robot.

Second, due to the higher number of hyperparameters, the optimization can take longer to finish compared to SKGP or SOGP. Fortunately, these are problems for MOGP in general [11], [12] and many solutions have been put forth to address them [26], [27]. Hence, the challenge is to find the right techniques and apply them to our intended use case, which is part of future work.

B. Real-world Dataset

In this section, the proposed system is evaluated on the Jura dataset [28], which includes a prediction set (259 sampling locations) and a test set (100 validation locations). The data include the concentration of seven heavy metals, corresponding coordinates, land use type, and rock type. In this work, we envision an environmental monitoring scenario where an autonomous robot can go to any sampling location to collect measurements for two elements - cadmium (Cd) and lead (Pb) - but cannot access the validation locations. Additionally, the concentration data for copper (Cu) for all sampling locations are available thanks to previous geostatistical surveys. These three elements were chosen since among the seven metals, the strongest and weakest correlations were respectively observed between Cu-Pb and Cu-Cd pairs [28]. As such, using them might help to elucidate the effectiveness of the proposed system. With this setting, the problems are: 1) predicting the concentrations of Cd (QoI_1) and Pb (QoI_2) at 100 validation locations using on-site measurements and a-priori data for Cu (PQ_1), and 2) selecting the next best location to sample among the unvisited locations. The dimensions of the problems are: $\mathbf{x} \in \mathbb{R}^2$, $\mathbf{y} \in \mathbb{R}^3$, $q = 2$, $p = 1$, $m = 259$, $M = 100$.

Fig. 8 presents the distribution maps for the QoIs and PQ. All measurements at sampling locations are affected by Gaussian noise with a standard deviation of 10% of the true value. Since the measurement ranges for the QoIs vary significantly, we use the percent error (PE), computed as $\text{PE}_{i,j} = \frac{|\hat{y}_i(\mathbf{x}_j) - y_i(\mathbf{x}_j)|}{y_i(\mathbf{x}_j)} \cdot 100\%$, to compare errors between different quantities. Table III shows the numerical results, while Fig. (9) illustrates the percent error of different modeling techniques with the Jura dataset. Firstly, it is obvious that MKGP achieves the lowest mean PE for both QoIs, which demonstrates the superior performance of our method. Secondly, MKGP obtains the lowest standard deviation (σ) and the lowest maximum PE, which are both significantly better than what can be achieved by either SOGP or SKGP. Thirdly, while all methods recognize the weak correlation between Cd and Cu, with only minor variations in their estimated coefficients, MKGP identifies the strong correlation in the Pb-Cu pair with the highest confidence. Overall, it is clear that MKGP outperforms previous approaches in this benchmark.

Nonetheless, challenges remain when working with real-world datasets, particularly in handling anomalies. As shown in Fig. (8), the values at certain validation locations, such as (V1)–(V3), exhibit drastic changes compared to their surroundings. Consequently, regardless of the modeling approach used, the PEs at these points are significantly higher than at other locations. However, as seen in Fig. (9), at (V3) MKGP produces a noticeably more accurate estimate than the other methods. This improvement may be attributed to the presence of a nearby sampling location (S2), where the a-priori Cu concentration also increases sharply relative to its surroundings. By leveraging this data point and recognizing the strong correlation between Pb and Cu, MKGP is able to

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

predict the Pb value at (V3) more effectively compared to SOGP and SKGP. In contrast, no similarly correlated a-priori data exist for (V1) and (V2), making accurate predictions at these locations challenging for all methods.

Finally, Fig. (10) illustrates the performance and cost of various sample selection methods. All methods are paired with MKGP and begin at the same starting point (S1). The results of the random sampler are averaged over 10 runs. For simplicity, the Euclidean distance between sample locations is used as the travel cost $\Lambda(\cdot)$. In general, to achieve accuracy within 3% of the final mean PE, CS and RS need more than 200 samples, while it only takes 158 samples for MVAS. Furthermore, it is obvious that the proposed MVAS method provides the fastest rate of improvement within the first 100 samples. However, to cover all sampling locations, the CS, RS, and MVAS methods require 489.4 km, 530.1 km, and 659.6 km of travel distance, respectively. These results show that MVAS offers the most cost-effective sampling strategy compared to other approaches by prioritizing the most informative samples first and taking into account the associated travel costs. However, it might not be the best option if the goal is to simply collect all of the samples in the environment in the shortest amount of time.

V. Conclusion

In this article, we present a system that addresses the robotic problem of environmental mapping for unknown scalar fields with limited resources. Our method includes MKGP, a multi-kernel Gaussian processes module and MVAS, a multivariate intelligent adaptive sampling module. MKGP first assigns a different kernel for each task, then uses convolution and free-form parameterization to formulate the cross-covariance terms. MVAS will leverage the learned model from MKGP, then incorporate prior information and travel cost to find the next best sampling location. In simulations with synthetic and real-world datasets, we show that our system can model the QoIs more accurately than SOGP and SKGP, discover the hidden relationships between QoIs and PQ more precisely than the Pearson correlation coefficient, and plan the sampling route more efficiently than previous approaches. In the future, we aim to improve our system by addressing the scaling issue with appropriate techniques, making our MVAS method more robust against anomalies in real-world data and more adaptive to the dynamics of the environment by including the coefficients α and β in the optimization problem, and conducting field trials to evaluate the system in the real world.

References

[1] H. Lipson and S. Sukkarieh, "Robots may transform the way we produce and prepare food," *Nature Reviews Bioengineering*, vol. 1, no. 11, pp. 795–798, 2023.

[2] A. Botta, P. Cavallone, L. Baglieri, G. Colucci, L. Tagliavini, and G. Quaglia, "A review of robots, perception, and tasks in precision agriculture," *applied mechanics*, vol. 3, no. 3, pp. 830–854, 2022.

[3] A. Shamshirgaran and S. Carpin, "Reconstructing a spatial field with an autonomous robot under a budget constraint," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 8963–8970.

[4] H. Zhu, J. J. Chung, N. R. Lawrance, R. Siegwart, and J. Alonso-Mora, "Online Informative Path Planning for Active Information Gathering of a 3D Surface," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 1488–1494.

[5] A. Candela, S. Kodgule, K. Edelson, S. Vijayarangan, D. R. Thompson, E. N. Dobrea, and D. Wettergreen, "Planetary rover exploration combining remote and in situ measurements for active spectroscopic mapping," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 5986–5993.

[6] T. Homberger, L. Wellhausen, P. Fankhauser, and M. Hutter, "Support surface estimation for legged robots," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8470–8476.

[7] C. Williams, S. Klanke, S. Vijayakumar, and K. Chai, "Multi-task gaussian process learning of robot inverse dynamics," *Advances in neural information processing systems*, vol. 21, 2008.

[8] J. Li, Y. Cheng, and J. Chen, "Multi-task gaussian process classification-based collaborative map fusion using air-ground robotic system," in *2020 International Symposium on Autonomous Systems (ISAS)*. IEEE, 2020, pp. 174–179.

[9] N. Harrison, N. Wallace, and S. Sukkarieh, "Automated Testing of Spatially-Dependent Environmental Hypotheses through Active Transfer Learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

[10] D. Korkinof and Y. Demiris, "Multi-task and multi-kernel gaussian process dynamical systems," *Pattern Recognition*, vol. 66, pp. 190–201, 2017.

[11] A. Melkumyan and F. Ramos, "Multi-kernel gaussian processes," in *Twenty-second international joint conference on artificial intelligence*, 2011.

[12] R. Dürichen, M. A. Pimentel, L. Clifton, A. Schweikard, and D. A. Clifton, "Multitask gaussian processes for multivariate physiological time-series analysis," *IEEE Transactions on Biomedical Engineering*, vol. 62, no. 1, pp. 314–322, 2014.

[13] J. Lee, J. Feng, M. Humt, M. G. Müller, and R. Triebel, "Trust your robots! predictive uncertainty estimation of neural networks with sparse gaussian processes," in *Proceedings of the 5th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Faust, D. Hsu, and G. Neumann, Eds., vol. 164. PMLR, 08–11 Nov 2022, pp. 1168–1179.

[14] A. Melkumyan and F. T. Ramos, "A sparse covariance function for exact gaussian process inference in large datasets," in *Twenty-first international joint conference on artificial intelligence*, 2009.

[15] Y. Sung, Z. Chen, J. Das, P. Tokekar *et al.*, "A survey of decision-theoretic approaches for robotic environmental monitoring," *Foundations and Trends® in Robotics*, vol. 11, no. 4, pp. 225–315, 2023.

[16] Y. Kantaros, B. Schlotfeldt, N. Atanasov, and G. J. Pappas, "Sampling-based planning for non-myopic multi-robot information gathering," *Autonomous Robots*, vol. 45, no. 7, pp. 1029–1046, 2021.

[17] J. P. Kleijnen and E. Mehdad, "Multivariate versus univariate kriging metamodels for multi-response simulation models," *European Journal of Operational Research*, vol. 236, no. 2, pp. 573–582, 2014.

[18] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.

[19] E. V. Bonilla, K. Chai, and C. Williams, "Multi-task gaussian process prediction," *Advances in neural information processing systems*, vol. 20, 2007.

[20] M. A. Alvarez and N. D. Lawrence, "Computationally efficient convolved multiple output gaussian processes," *The Journal of Machine Learning Research*, vol. 12, pp. 1459–1500, 2011.

[21] G. Wynne, F.-X. Briol, and M. Girolami, "Convergence guarantees for gaussian process means with misspecified likelihoods and smoothness," *Journal of Machine Learning Research*, vol. 22, no. 123, pp. 1–40, 2021.

[22] J. D. Svenson and T. J. Santner, "Multiobjective optimization of expensive black-box functions via expected maximin improvement," *The Ohio State University, Columbus, Ohio*, vol. 32, 2010.

[23] C. Q. Lam, "Sequential adaptive designs in computer experiments for response surface model fit," Dissertation, The Ohio State University, 2008.

[24] "NSW Potassium Percentage (K%)," https://geonetwork.geoscience.nsw.gov.au/geonetwork/srv/eng/catalog.search#/metadata/341e919f-8e36-47c8-8947-ce5e2b91287f, accessed: 2024-09-05.

[25] D. P. Kroese and Z. I. Botev, "Spatial process simulation," in *Stochastic geometry, spatial statistics and random fields: Models and algorithms*. Springer, 2014, pp. 369–404.

[26] M. Nabati, S. A. Ghorashi, and R. Shahbazian, "JGPR: a computationally efficient multi-target Gaussian process regression algorithm," *Machine Learning*, vol. 111, no. 6, p. 1987–2010, 2022.

[27] H. Liu, J. Ding, X. Xie, X. Jiang, Y. Zhao, and X. Wang, "Scalable multi-task gaussian processes with neural embedding of coregionalization," *Knowledge-Based Systems*, vol. 247, p. 108775, 2022.

[28] P. Goovaerts, *Geostatistics for natural resources evaluation*. Oxford University Press, USA, 1997.